Enabling the Next Generation of Photonics through Design IP Reuse

Matthew Streshinsky
Enosemi, Inc.
San Jose, CA, U.S.A.
matt.streshinsky@enosemi.com

Ari Novack
Enosemi, Inc.
San Jose, CA, U.S.A.
ari.novack@enosemi.com

Shahab Ardalan
Enosemi, Inc.
San Jose, CA, U.S.A.
shahab.ardalan@enosemi.com

***Abstract*—With the emergence of several open-market silicon photonics foundries, it is now possible to create a rich design IP ecosystem for co-designed electronics and photonics much like what exists for CMOS. Design reuse will accelerate time-to-market and enable significantly enhanced photonic system complexity.**

Keywords—Silicon photonics, optical communication, very large scale integration (VLSI) photonics

I. Introduction

Historically, many of the non-academic silicon photonics chip manufacturers utilized closed or proprietary processes [1-5]. Traditionally, the target application for silicon photonic devices built in these processes were transceivers for 10 gigabaud/lane to 50 gigabaud/lane direct detect or coherent communication systems with up to 4 lanes per chip. In these applications, silicon-based optical devices typically entered the market after "traditional-technology-based" device implementations such as EMLs were demonstrated [6]. Silicon photonics directly tackled performance, power, cost, and volume manufacturability at these data rates. From vendor to vendor, the silicon photonics devices that were deployed differed from each other in terms of their respective semiconductor manufacturing process flow.

Looking forward, there are a tremendous quantity of new opportunities to develop significantly more complex devices for switching and novel compute systems [7]. The degree of integration present in the latest generation of silicon photonics chips are now in the realm of "very large-scale integration" (VLSI), based on the number of devices per chip [8]. As the target applications increase in complexity, the industry will face an increasing need for design reuse, otherwise every entrant will be forced to re-engineer the same fundamental building blocks repeatedly. While the lack of design reuse may have enabled competitive market edges for typical transceiver players, it will encumber the broader silicon photonics industry's ability to innovate.

In recent years, several platforms and PDK offerings have emerged, offering open market access to high-performance photonic devices in state-of-the-art manufacturing facilities [9-14]. These platforms are developed by leading semiconductor organizations such as GlobalFoundries, Tower Semi, IMEC, Taiwan Semiconductor Manufacturing Corporation, and others. Access to stable processes will enable the development of reusable devices, circuits, systems, and chiplets. Further, this will enable a new wave of products to launch with the same confidence and process control to which the electronics industry is accustomed. Design reuse will enable the industry to rapidly advance to and beyond VLSI photonics and enable the next generation of photonics. In this manuscript, we aim to motivate a design IP hierarchy of photonics with co-designed electronics and present an example use case by way of a 16-channel monolithically integrated transmitter and receiver built with reusable components.

II. Co-designed electronic and photonic design IP

Silicon photonic design IP, much like CMOS design IP, can exist in a wide variety of complexities. Built on top of the foundry's process design kit (PDK) are essential device designs. Individual devices can be passive elements such as multi-mode interferometers or active devices such as ring modulators or Mach-Zehnder modulators (which in-turn already re-use one or more passive devices and PDK elements). Design reuse does not have to end at this level of complexity, though. It is possible to assemble those devices into larger systems: an IQ modulator, a wavelength demultiplexer, or a coherent receiver, for example. Further, in combination with co-designed electronic design IP, it should be possible to instantiate a fully functional and stabilized Mach-Zehnder interferometer (MZI) with built-in feedback control. With one such MZI block implemented, it is possible to cascade many such IP blocks into a full photonic tensor processor [15]. Now, such a processor core can in turn be reused and implemented in-line with other photonic systems.

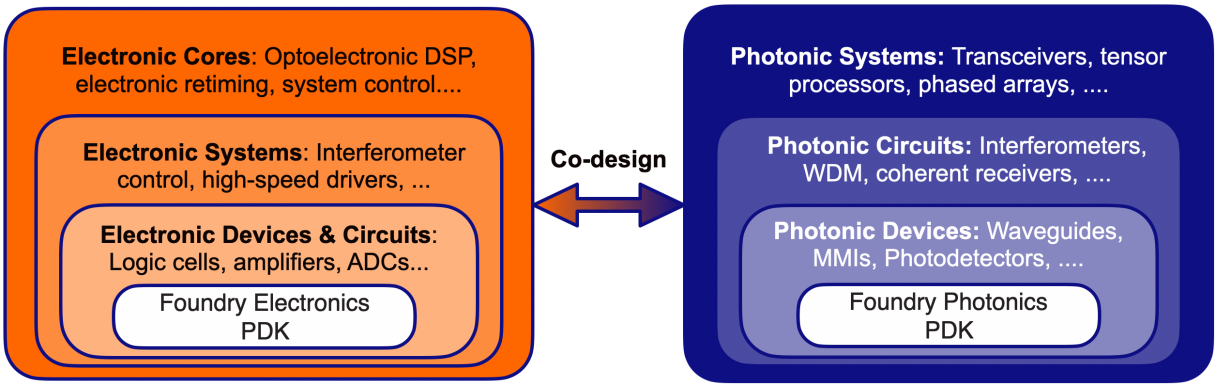

Fig. 1. Electronic-photonic design hierarchy. VLSI photonics will require the close integration and design of both disciplines.

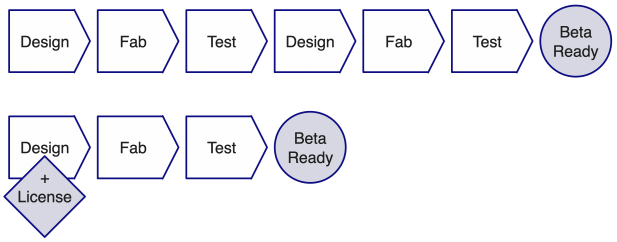

Fig. 2. Design IP reuse accelerates time-to-market by enabling first-silicon-success with verified designs require the close integration and

Whether the electronics are monolithically integrated or on separate die, silicon photonic chips will almost always require electronic interfaces. While post-fabrication trimming technology has emerged for certain devices requiring precise optical phase control, truly complex photonic systems will necessarily require electronics. The electronics may be used for both feedback control systems to operate across process, voltage, and temperature corners; as well as create interesting dynamic systems.

As the systems grow in complexity, co-design becomes increasingly necessary. Technology, process, and performance become coupled with outcomes. VLSI photonics requires tight coupling between the various photonic elements with the electronic I/O, electronic control, and electronic processing. For example, four channels of Mach-Zehnder interferometer controllers can be simply implemented with commercial off-the-shelf ADCs, DACs, Op-Amps, and microcontrollers in a small form factor pluggable, evidenced by myriad transceivers on sale today. One thousand Mach-Zehnder interferometers, on the other hand, stand no chance of achieving the same form factor without integration in much the same way it is possible but usually inadvisable to build a computer system from discretely packaged DIP logic gates.

III. Enabling the Next Generation of Photonics

Co-design of electronics and photonics is motivated by more than just integration and performance advantages. Reusing circuits purposefully engineered as a complete solution radically reduces risk. Returning to the Mach-Zehnder interferometer example, the ability to re-use a block that is designed for reuse means the control algorithms are validated, the layout will work, the system can be tested in the lab prior to tapeout, and the specifications for this component are already known. The probability that a redesign is required because the control system is faulty is significantly reduced compared to if it were being designed from scratch.

In order to confidently reuse a piece of design IP, it is necessary to trust that the design IP will work. Thus, design IP becomes valuable when it has been validated. Having access to validated design IP means that highly complex circuit designs can work the first time, so called "first silicon success." See fig. 2 for an illustration of this concept. Design cycles that require only one tapeout now means that both riskier designs can be pursued with the same R&D effort, and that the overall time to market has been accelerated. For any compounding process, a factor of two change in the exponent will yield large differences in outcomes. This can be leveraged by any R&D team to either develop de-risked designs or to develop radically more complicated designs with the same overall design risk at tapeout. Thus, validated design IP will accelerate the entire silicon photonics industry and fundamentally enable new opportunities.

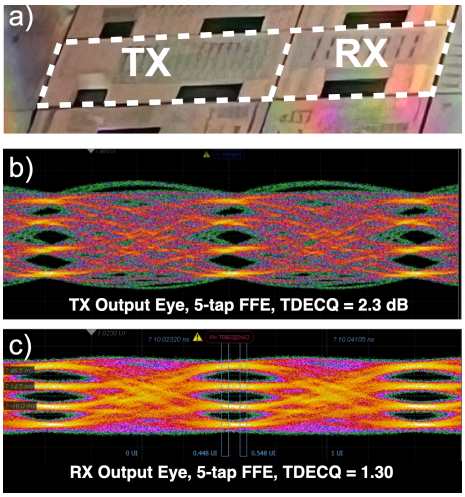

Fig. 3. a) Example 16-channel TX and RX chip. b) TX test chip output eye at 112 Gbps PAM4. c) RX test chip output eye at 112 Gbps PAM4.

Standardized processes and a design IP ecosystem enable the silicon photonics industry to move away from the design philosophy where a design team interacts with a foundry by asking, “what process changes can we make to improve the design?” Instead, the design team can ask, “how can complexity improve the design?” While it isn’t guaranteed that complexity will solve all design problems, it will enable a new axis of optimization that did not exist before. Differentiation by complexity, rather than by process and materials, is a relatively uncommon concept in the photonics industry. Of course, the foundries will continue to innovate on the process and with novel materials integration, but this innovation can occur in lockstep with photonic system design complexity similarly to how CMOS technology has advanced to its present state.

Finally, there is a positive feedback loop that may emerge. Standardized processes, in combination with a design IP ecosystem, in combination with a plethora of new applications and systems, leads to a scenario where there is increased investment in the aforementioned areas, not to mention the economies of scale that come with running a manufacturing process in higher volumes. Further, the need for design IP will necessitate more design IP, which can in turn enable new opportunities.

IV. Example Monolithic CMOS-Silicon Photonics IP

Here, we demonstrate an example of design reuse for optical I/O applications with a 16-channel monolithically integrated silicon photonics transmitter and receiver. Additional details of these devices are available in ref. [16, 17]. Inside the transmitter there are drivers, traveling wave Mach-Zehnder modulators, feedback control systems, and digital interfaces integrated monolithically on-chip. Implementing sixteen copies of this transmit (TX) and receive (RX) macro can be simply achieved by placing sixteen copies side-by-side in accordance with the design IP and foundry PDK's design rules. In doing so, the full multi-channel transmit chip is achieved, with a die image shown in Fig. 3(a). In this example, we demonstrate that both the transmitter and receiver devices operate at 112 Gbps PAM4, with eye diagrams shown in Fig. 3(b, c).

Importantly, each of the sub-components inside the TX and RX macros are themselves individually useful. For

example, inside the TX macro is a Mach-Zehnder interferometer automatic bias control system. The control system can automatically stabilize the Mach-Zehnder modulators at the optical quadrature bias point settable via SPI bus. The system operates by comparing photocurrent readouts tapped from in-line waveguides at both outputs of the Mach-Zehnder 2x2MMI combiner. The control system goes on to compare the resulting photocurrents to determine whether to increment or decrement the power dissipated on one of the thermal tuners in the modulator. As part of that control system, low speed transimpedance amplifiers, ADCs, DACs, and low dropout regulators are developed and implemented.

These various stand-alone mixed-signal devices or sub-systems, such as an optical power monitor, can be reused throughout a chip layout or in a chip design within nearly any other application space. Using these devices enables, for example, the aforementioned photonic tensor processor to be implemented as a single chip. DACs and LDOs could drive the series of tunable coupler thermal phase shifters, TIAs and DACs could receive the output, and new integrated digital logic could control it all. Crucially, the scale of the system, now as a single chip, could be increased since the number of analog inputs and outputs are no longer constrained by electronic packaging limitations.

Access to these devices obviates the need for another silicon photonic effort to re-design, re-tapeout, and re-test the same structures, and do it with the knowledge that this design will work and meet their respective needs. Core to this enablement is the requirement to not modify the fundamental process. If the manufacturing process changes, then the designs can no longer be guaranteed to the same degree of certainty.

V. Conclusion

Validated photonic devices and electronic-photonic circuits accelerates time-to-market and enables significantly more complex designs than what was previously possible within a certain risk profile. Open market standardized manufacturing processes enables this ecosystem to exist. Once the combination of a design IP ecosystem is present, truly novel devices and designs can emerge.

We presented a first step to enable this ecosystem by providing a few of the fundamental building blocks for complex integrated system: high-speed linear drivers and TIAs, DAC and ADCs, a variety of analog circuitry, and digital control systems to stabilize optically phase-sensitive paths. We aim for the broader silicon photonics industry to begin using these devices to build far more complicated systems than four-channel transmitter and receiver arrays.

While historically, silicon photonic products differentiated among one another by manufacturing process, the next generation of devices can differentiate through complexity. What exactly is the next generation of photonics? The first step to finding out is to stop repeatedly reinventing the same photonic circuits and rather use them to start building.

Acknowledgment

The authors thank Luminous Computing for providing the original 112 Gbps PAM4 test data and silicon photonic assemblies, as well as the GlobalFoundries team for design and manufacturing support.

References

[1] C. Doerr and L. Chen, "Silicon Photonics in Optical Coherent Systems," Proceedings of the IEEE, vol. 106, no. 12, pp. 2291-2301, 2018.

[2] Gregory Fish, "High Bandwidth Transceivers Using Heterogeneous Integration on Silicon," OFC 2015, paper W3H.2.

[3] A. Novack, et. al. “A silicon photonic transceiver and hybrid tunable laser for 64 Gbaud coherent communication,” OFC 2018, paper Th4D.4.

[4] A. Mekis, et al., "A Grating-Coupler-Enabled CMOS Photonics Platform," IEEE IEEE J. Quantum Electron, vol. 17, no. 3, pp. 597-608, 2011.

[5] R. Jones et al., "Heterogeneously Integrated InP/Silicon Photonics: Fabricating Fully Functional Transceivers", IEEE Nanotechnology Mag., vol. 13, no. 2, pp. 17-26, 2019.

[6] M. Bohn, P. Magill, M. Hochberg, D. Scordo, A. Novack, and M. Streshinsky, "Next-Generation Silicon Photonic Interconnect Solutions," OFC 2019, paper M3J.3.

[7] S. Shekhar, et al., "Silicon Photonics - Roadmapping the Next Generation," arXiv:2305.15820v1, 2023.

[8] N. Margalit, et al., "Perspective on the future of silicon photonics and electronics," Applied Physics Letters, vol. 118, no. 22, 220501 (2021).

[9] M. Rakowski, et al., "45nm CMOS - Silicon Photonics Monolithic Technology (45CLO) for next-generation, low power and high speed optical interconnects," OFC 2020, paper T3H.3.

[10] Tower Semi, https://towersemi.com/technology/rf-and-hpa/silicon-photonics-rf/

[11] AMF, “Advanced Micro Foundry,” https://www.advmf.com/service/

[12] VTT, “VTT,” https://www.vttresearch.com/en/ourservices/photonics

[13] IHP, “IHP,” https://www.ihp-solutions.com/

[14] AIM Photonics, “AIM Photonics (SUNY),”http://www.aimphotonics.com/pdk

[15] S. Bandyopadhyay, et al., "A Photonic Deep Neural Network Processor on a Single Chip with Optically Accelerated Training," CLEO 2023, paper SM2P.2.

[16] T. Baehr-Jones, et al., "Monolithically integrated 112 Gbps PAM4 optical transmitter and receiver in a 45 nm CMOS-silicon photonics process," Optics Express, vol. 31, no. 15, pp. 24926-24938, 2023.

[17] A. Darbinain, et al.,, “Monolithic Mach-Zehnder Modulator Bias Controls in a 45 nm Silicon Photonics Process,” ECOC 2023, in press.